arXiv:1510.02354v1 [stat.ML] 8 Oct 2015

The Knowledge Gradient with Logistic Belief Models for Binary Classification

Yingfei Wang * Chu Wang † Warren Powell ‡

Abstract

We consider sequential decision making problems for binary classification scenario in which the learner takes an active role in repeatedly selecting samples from the action pool and receives the binary label of the selected alternatives. Our problem is motivated by applications where observations are time consuming and/or expensive, resulting in small samples. The goal is to identify the best alternative with the highest response. We use Bayesian logistic regression to predict the response of each alternative. By formulating the problem as a Markov decision process, we develop a knowledge-gradient type policy to guide the experiment by maximizing the expected value of information of labeling each alternative and provide a finite-time analysis on the estimated error. Experiments on benchmark UCI datasets demonstrate the effectiveness of the proposed method.

1 Introduction

There are many classification problems where observations are time consuming and/or expensive. One example arises in health care analytics, where physicians have to make medical decisions (e.g. a course of drugs, surgery, and expensive tests). Assume that a doctor faces a discrete set of medical choices, and that we can characterize an outcome as a success (patient does not need to return for more treatment) or a failure (patient does need followup care such as repeated operations). We encounter two challenges. First, there are very few patients with the same characteristics, creating few opportunities to test a treatment. Second, testing a medical decision may require several weeks to determine the outcome. This creates a situation where experiments (evaluating a treatment decision) are time consuming and expensive, requiring that we learn from our decisions as quickly as possible.

The challenge of deciding which medical decisions to evaluate can be modeled mathematically as a sequential decision making problem with binary outcomes. In this setting, we have a budget of measurements that we allocate sequentially to medical decisions so that when we finish our study, we have collected information to maximize our ability to identify the best medical decision with the highest response (probability of success). Scientists can draw on an extensive body of literature on the classic design of experiments [1, 2, 3] whose goal is to decide what observations to make when fitting a function. Yet in the laboratory settings considered in this paper, the decisions need to be guided by a well-defined utility function (that is, identify the best alternative with the highest probability of success). This

*Department of Computer Science, Princeton University, yingfei@cs.princeton.edu
†Program in Applied and Computational Mathematics, Princeton University, chuw@math.princeton.edu
‡Department of Operations Research and Financial Engineering, Princeton University, powell@princeton.edu

problem also relates to active learning [4, 5, 6, 7] in several aspects. In terms of active learning scenarios, our model is most similar to membership query synthesis where the learner may request labels for any unlabeled instance in the input space to learn a classifier that accurately predicts the labels of new examples. By contrast, our goal is to maximize a utility function such as the success of a treatment. Also, it is typical in active learning not to query a label more than once, whereas we have to live with noisy outcomes, requiring that we sample the same label multiple times. Moreover, the expense of labeling each alternative sharpens the conflicts of learning the prediction and finding the best alternative. Another similar sequential decision making setting is multi-armed bandit problems (e.g. [8, 9]). Our work will initially focus on offline settings such as laboratory experiments or medical trials, but the knowledge gradient for offline learning extends easily to online settings [10].

There is a literature studying sequential decision problems to maximize a utility function (e.g., [11, 12, 13]). We are particularly interested in a policy that is called the knowledge gradient (KG) that maximizes the expected value of information. After its first appearance for ranking and selection problems [14], KG has been extended to various other belief models (e.g. [15, 16, 10, 17]). Yet there is no KG variant designed for binary classification with parametric models. In this paper, we extend the KG policy to the setting of classification problems under a logistic belief model which introduces the computational challenge of working with nonlinear models.

This paper is organized as follows. We first rigorously establish a sound mathematic model for the problem of sequentially maximizing the response under binary outcome in Section 2. We then develop a recursive Bayesian logistic regression procedure to predict the response of each alternative and further formulate the problem as a Markov decision process. In Section 3, we design a knowledge-gradient type policy under a logistic belief model to guide the experiment and provide a finite-time analysis on the estimated error. This is different from the PAC (passive) learning bound which relies on the i.i.d. assumption of the examples. Experiments are demonstrated in Section 4.

2 Problem formulation

In this section, we state a formal model for our response maximization problem, including transition and objective functions. We then formulate the problem as a Markov decision process.

2.1 The mathematical model

We assume that we have a finite set of alternatives $\boldsymbol{x} \in \mathcal{X} = \{\boldsymbol{x}_1, \ldots, \boldsymbol{x}_M\}$. The observation of measuring each $\boldsymbol{x}$ is a binary outcome $y \in \{-1, +1\}$ with some unknown probability $\Pr[y = +1|\boldsymbol{x}]$. Under a limited budget N, our goal is to choose the measurement policy $(\boldsymbol{x}^1, \ldots, \boldsymbol{x}^N)$ and implementation decision $\boldsymbol{x}^{N+1}$ that maximizes $\Pr(y = +1|\boldsymbol{x}^{N+1})$. We assume a parametric model where each $\boldsymbol{x}$ is a d-dimensional vector and the probability of an example $\boldsymbol{x}$ belonging to class $+1$ is given by a nonlinear transformation of an underlying linear function of $\boldsymbol{x}$ with a weight vector $\boldsymbol{w}$:

$$\Pr(y = +1|\boldsymbol{x}, \boldsymbol{w}) = \sigma(\boldsymbol{w}^T \boldsymbol{x}),$$

with the sigmoid function $\sigma(a)$ chosen as the logistic function $\sigma(a) = \frac{1}{1+\exp(-a)}$.

We assume a Bayesian setting in which we have a multivariate prior distribution for the unknown parameter vector $\boldsymbol{w}$. At iteration n, we choose an alternative $\boldsymbol{x}^n$ to measure and observe a binary outcome y^n assuming labels are generated independently given $\boldsymbol{w}$.

Each alternative can be measured more than once with potentially different outcomes. Let $\mathcal{D}^n = \{(\boldsymbol{x}^i, y^i)\}_{i=1}^n$ denote the previous measured data set for any $n = 1, \ldots, N$. Define the filtration $(\mathcal{F}^n)_{n=0}^N$ by letting $\mathcal{F}^n$ be the sigma-algebra generated by $\boldsymbol{x}^1, y^1, \ldots, \boldsymbol{x}^n, y^n$. We use $\mathcal{F}^n$ and $\mathcal{D}^n$ interchangeably. Measurement and implementation decisions $\boldsymbol{x}^{n+1}$ are restricted to be $\mathcal{F}^n$-measurable so that decisions may only depend on measurements made in the past. We use Bayes' rule to form a sequence of posterior predictive distributions $\Pr(\boldsymbol{w}|\mathcal{D}^n)$ for $\boldsymbol{w}$ from the prior and the previous measurements.

The next lemma states the equivalence of using true probabilities and sample estimates when evaluating a policy, where Π is the set of policies. The proof is left in the supplementary material.

Lemma 1. *Let $\pi \in \Pi$ be a policy, and $\boldsymbol{x}^\pi = \arg\max_{\boldsymbol{x}} Pr[y = +1|\boldsymbol{x}, \mathcal{D}^N]$ be the implementation decision after the budget N is exhausted. Then*

$$\mathbb{E}[Pr(y = +1|\boldsymbol{x}^\pi, \boldsymbol{w})] = \mathbb{E}[\max_{\boldsymbol{x}} Pr(y = +1|\boldsymbol{x}, \mathcal{D}^N)],$$

where the expectation is taking over the prior distribution of $\boldsymbol{w}$.

By denoting $\mathcal{X}^I$ as an implementation policy for selecting an alternative after the measurement budget is exhausted, then $\mathcal{X}^I$ is a mapping from the history $\mathcal{D}^N$ to an alternative $\mathcal{X}^I(\mathcal{D}^N)$. Then as a corollary of Lemma 1, we have [13]

$$\max_{\mathcal{X}^I} \mathbb{E}\big[\Pr\big(y = +1|\mathcal{X}(\mathcal{D}^N)\big)\big] = \max_{\boldsymbol{x}} \Pr(y = +1|\boldsymbol{x}, \mathcal{D}^N).$$

In other words, the optimal decision at time N is to go with our final set of beliefs. By the equivalence of using true probabilities and sample estimates when evaluating a policy, while we want to learn the unknown true value $\max_{\boldsymbol{x}} \Pr(y = +1|\boldsymbol{x})$, we may write our problem's objective as

$$\max_{\pi \in \Pi} \mathbb{E}^\pi [\max_{\boldsymbol{x}} \Pr(y = +1|\boldsymbol{x}, \mathcal{D}^N)]. \tag{1}$$

2.2 From logistic regression to Markov decision process formulation

Logistic regression is widely used in machine learning for binary classification [18]. Given a training set $\mathcal{D} = \{(\boldsymbol{x}_i, y_i)\}_{i=1}^n$ with $\boldsymbol{x}_i$ a d-dimensional vector and $y_i \in \{-1, +1\}$, with the assumption that training labels are generated independently given $\boldsymbol{w}$, the likelihood $\Pr(\mathcal{D}|\boldsymbol{w})$ is defined as $\Pr(\mathcal{D}|\boldsymbol{w}) = \prod_{i=1}^n \sigma(y_i \cdot \boldsymbol{w}^T \boldsymbol{x}_i)$. In frequentists' interpretation, the weight vector $\boldsymbol{w}$ is found by maximizing the likelihood of the training data $\Pr(\mathcal{D}|\boldsymbol{w})$. l_2-regularization has been used to avoid over-fitting with the estimate of the weight vector $\boldsymbol{w}$ given by:

$$\min_{\boldsymbol{w}} \frac{\lambda}{2}\|\boldsymbol{w}\|^2 + \sum_{i=1}^n \log(1 + \exp(-y_i \boldsymbol{w}^T \boldsymbol{x}_i)). \tag{2}$$

2.2.1 Bayesian setup

Exact Bayesian inference for logistic regression is intractable since the evaluation of the posterior distribution comprises a product of logistic sigmoid functions and the integral in the normalization constant is intractable as well. With a Gaussian prior on the weight vector, the Laplace approximation can be obtained by finding the mode of the posterior distribution and then fitting a Gaussian distribution centered at that mode (see Chapter 4.5 of [19]). Specifically, suppose we begin with a Gaussian prior $\Pr(\boldsymbol{w}) = \mathcal{N}(\boldsymbol{w}|\boldsymbol{m}, \boldsymbol{\Sigma})$, and we

wish to approximate the posterior $\Pr(\boldsymbol{w}|\mathcal{D}) \propto \Pr(\mathcal{D}|\boldsymbol{w})\Pr(\boldsymbol{w})$. Define the logarithm of the unnormalized posterior distribution

$$\begin{aligned}\Psi(\boldsymbol{w}|\boldsymbol{m}, \boldsymbol{\Sigma}, \mathcal{D}) &= \log \Pr(\mathcal{D}|\boldsymbol{w}) + \log \Pr(\boldsymbol{w}) \\ &= -\frac{1}{2}(\boldsymbol{w}-\boldsymbol{m})^T\boldsymbol{\Sigma}^{-1}(\boldsymbol{w}-\boldsymbol{m}) - \sum_{i=1}^{n}\log(1+\exp(-y_i\boldsymbol{w}^T\boldsymbol{x}_i)). \qquad (3)\end{aligned}$$

The Laplace approximation is based on a Taylor expansion to Ψ around its MAP (maximum a posteriori) solution $\hat{\boldsymbol{w}} = \arg\max_{\boldsymbol{w}} \Psi(\boldsymbol{w})$, which defines the mean of the Gaussian. The covariance is then given by the Hessian of the negative log posterior evaluated at $\hat{\boldsymbol{w}}$, which takes the form

$$(\boldsymbol{\Sigma}')^{-1} = -\nabla^2\Psi(\boldsymbol{w})|_{\boldsymbol{w}=\hat{\boldsymbol{w}}} = \boldsymbol{\Sigma}^{-1} + \sum_{i=1}^{n} p_i(1-p_i)\boldsymbol{x}_i\boldsymbol{x}_i^T, \qquad (4)$$

where $p_i = \sigma(\hat{\boldsymbol{w}}^T\boldsymbol{x}_i)$. The Laplace approximation results in a normal approximation to the posterior

$$\Pr(\boldsymbol{w}|\mathcal{D}) \approx \mathcal{N}(\hat{\boldsymbol{w}}, \boldsymbol{\Sigma}'). \qquad (5)$$

By substituting an independent normal prior with q_i^{-1} as the diagonal element of diagonal covariance matrix $\boldsymbol{\Sigma}$, the Laplace approximation to the posterior distribution of each weight w_j reduces to $\Pr(w_j|\mathcal{D}) \approx \mathcal{N}(\hat{w}_j, q_j^{-1})$. Note here if $q_j = \lambda, m_i = 0$, the solution of Eq. (2) is the same as the MAP solution of (3). So an l_2-regularized logistic regression can be interpreted as a Bayesian model with a Gaussian prior on the weights with standard deviation $1/\sqrt{\lambda}$.

2.2.2 Recursive Bayesian logistic update

Our state space is the space of all possible predictive distributions for $\boldsymbol{w}$. Starting from a Gaussian prior $\mathcal{N}(\boldsymbol{w}|\boldsymbol{m}^0, \boldsymbol{\Sigma}^0)$, after the first n observed data, the Laplace approximated posterior distribution is $\Pr(\boldsymbol{w}|\mathcal{D}^n) \approx \mathcal{N}(\boldsymbol{w}|\boldsymbol{m}^n, \boldsymbol{\Sigma}^n)$ according to (5). We formally define the state space $\mathcal{S}$ to be the cross-product of $\mathbb{R}^d$ and the space of positive semidefinite matrices. At each time n, our state of knowledge is thus $S^n = (\boldsymbol{m}^n, \boldsymbol{\Sigma}^n)$. Since retraining the logistic model using all the previous data after each new data comes in to update from S^n to S^{n+1} by obtaining the MAP solution of (3) or even with diagonal covariance with constant diagonal elements is clumsy, the Bayesian logistic regression can be extended to leverage for recursive model updates after each of the training data.

To be more specific, after the first training data, the Laplace approximated posterior is $\mathcal{N}(\boldsymbol{w}|\boldsymbol{m}^1, \boldsymbol{\Sigma}^1)$. This serves as a prior on the weights to update the model when the next training data becomes available. In this recursive way of model updating, previously measured data need not be stored or used for retraining the logistic model. For the rest of this paper, we focus on independent normal priors (with $\boldsymbol{\Sigma} = \lambda^{-1}\boldsymbol{I}$, where $\boldsymbol{I}$ is the identity matrix), which is equivalent to l_2-regularized logistic regression, which also offers greater computational efficiency. All the results can be easily generalized to the correlative normal case. By setting the batch size $n = 1$ and $\Sigma = \lambda^{-1}\boldsymbol{I}$ in Eq. (3) and (4), we have the recursive Bayesian logistic regression as in Algorithm 1.

Since $\Psi(\boldsymbol{w}|\boldsymbol{m}, \boldsymbol{\Sigma}, \mathcal{D})$ is convex in $\boldsymbol{w}$, we can tap a wide range of convex optimization algorithms including gradient search, conjugate gradient, an BFGS method (see [20] for details). Yet when setting $n = 1$ and $\Sigma = \lambda^{-1}\boldsymbol{I}$ in Eq. (3), a stable and efficient algorithm for solving $\arg\max \Psi(\boldsymbol{w}) = -\frac{1}{2}\sum_{j=1}^{d} q_i(w_i - m_i)^2 - \log(1+\exp(-y\boldsymbol{w}^T\boldsymbol{x}))$ can be obtained

Algorithm 1: Recursive Bayesian Logistic Regression

input : Regularization parameter $\lambda > 0$
$m_j = 0$, $q_j = \lambda$. (Each weight w_j has an independent prior $\mathcal{N}(m_j, q_j^{-1})$)
for $t = 1$ *to* T **do**
  Get a new point $(\boldsymbol{x}, y)$.
  Find $\hat{\boldsymbol{w}}$ as the maximizer of (3): $-\frac{1}{2}\sum_{j=1}^{d} q_i(w_i - m_i)^2 - \log(1 + \exp(-y\boldsymbol{w}^T\boldsymbol{x}))$.
  $m_j = \hat{w}_j$
  Update q_i according to (4): $q_j \leftarrow q_j + \sigma(\hat{\boldsymbol{w}}^T\boldsymbol{x})(1 - \sigma(\hat{\boldsymbol{w}}^T\boldsymbol{x})x_j^2$.
end

as follows. First we calculate

$$\frac{\partial F}{\partial w_i} = -q_i(w_i - m_i) + \frac{yx_i \exp(-y\mathbf{w}^T\mathbf{x})}{1 + \exp(-y\mathbf{w}^T\mathbf{x})}.$$

By setting $\partial F/\partial w_i = 0$ for all i, then by denoting $(1 + \exp(y\mathbf{w}^T\mathbf{x}))^{-1}$ as p, we have

$$q_i(w_i - m_i) = ypx_i, \qquad i = 1, 2, \ldots, d,$$

and thus $w_i = m_i + yp\frac{x_i}{q_i}$. Plugging in these equalities into the definition of p, we have

$$\frac{1}{p} = 1 + \exp\Big(y\sum_{i=1}^{d}(m_i + yp\frac{x_i}{q_i})x_i\Big) = 1 + \exp(y\mathbf{m}^T\mathbf{x})\exp\Big(y^2p\sum_{i=1}^{d}\frac{x_i^2}{q_i}\Big).$$

The left hand side decreases from infinity to 1 and the right hand side increases from 1 when p goes from 0 to 1, therefore the solution exists and is unique in $[0, 1]$. By reducing a d-dimensional problem to a 1-dimensional one, the simple bisection method is good enough.

2.3 Markov decision process formulation

Our learning problem is a dynamic program that can be formulated as a Markov decision process. By using diagonal covariance matrices, the state space degenerates to $\mathcal{S} := \mathbb{R}^d \times (0, \infty]^d$ and it consists of points $s = (\boldsymbol{m}, \boldsymbol{q})$, where m_i, q_i are the mean and the precision of a normal distribution. We next define the transition function based on the recursive Bayesian logistic regression.

Definition 1. *The transition function* T*:* $\mathcal{S} \times \mathcal{X} \times \{-1, 1\}$ *is defined as*

$$T\Big((\boldsymbol{m}, \boldsymbol{q}), \boldsymbol{x}, y\Big) = \Big(\Big(\hat{\boldsymbol{w}}(\boldsymbol{m}), \boldsymbol{q} + p(1-p)\boldsymbol{diag}(\boldsymbol{x}\boldsymbol{x}^T)\Big)\Big),$$

where $\hat{\boldsymbol{w}}(\boldsymbol{m}) = \arg\min_{\boldsymbol{w}} \Psi(\boldsymbol{w}|\boldsymbol{m}, \boldsymbol{q})$*,* $p = \sigma(\hat{\boldsymbol{w}}^T\boldsymbol{x})$ *and* $\boldsymbol{diag}(\boldsymbol{x}\boldsymbol{x}^T)$ *is a column vector containing the diagonal elements of* $\boldsymbol{x}\boldsymbol{x}^T$*, so that* $S^{n+1} = T(S^n, \boldsymbol{x}, y)$*.*

In a dynamic program, the value function is defined as the value of the optimal policy given a particular state S^n at time n, and may also be determined recursively through Bellman's equation. If the value function can be computed efficiently, the optimal policy may then also be computed from it. The value function $V^n : \mathcal{S} \mapsto \mathbb{R}$ at time $n = 1, \ldots, N + 1$ is given by (1) as

$$V^n(s) := \max_{\pi} \mathbb{E}^{\pi}[\max_{\boldsymbol{x}} \Pr(y = +1|\boldsymbol{x}, \mathcal{F}^N)|S^n = s].$$

By noting that $\max_{\boldsymbol{x}} \Pr(y=+1|\boldsymbol{x}, \mathcal{F}^N)$ is $\mathcal{F}^N$-measurable and thus the expectation does not depend on policy π, the terminal value function V^N can be computed directly as

$$V^{N+1}(s) = \max_{\boldsymbol{x}} \Pr(y=+1|\boldsymbol{x}, s), \forall s \in \mathcal{S}.$$

The value function at times $n = 1, \ldots, N$, V^n, is given recursively by

$$V^n(s) = \max_{\boldsymbol{x}} \mathbb{E}[V^{n+1}(T(s, \boldsymbol{x}, y))], s \in \mathcal{S}.$$

3 Knowledge gradient policy for logistic belief model

Since the "curse of dimensionality" makes direct computation of the value function intractable, computationally efficient approximate policies need to be considered. A computationally attractive policy for ranking and selection problems is the knowledge gradient (KG), which is a stationary policy that at the nth iteration chooses its $(n+1)$th measurement to maximize the single-period expected increase in value [14]. It enjoys some nice properties, including myopic and asymptotic optimality. After its first appearance, KG has been extended to various belief models (e.g. [15, 16, 10, 17]) for offline learning, and an immediate extension to online learning problems [10]. Yet there is no KG variant designed for binary classification with parametric models, primarily because of the complexity of dealing with nonlinear belief models. In what follows, we extend the KG policy to the setting of classification problems under a logistic belief model.

Definition 2. *The knowledge gradient of measuring an alternative $\boldsymbol{x}$ while in state s is*

$$\nu_{\boldsymbol{x}}^{KG}(s) := \mathbb{E}[V^{N+1}(T(s, \boldsymbol{x}, y)) - V^{N+1}(s)]. \tag{6}$$

Since the label for alternative $\boldsymbol{x}$ is not known at the time of selection, the expectation is computed conditional on the current model specified by $s = (\boldsymbol{m}, \boldsymbol{q})$. Specifically, given a state $s = (\boldsymbol{m}, \boldsymbol{q})$, the label y for an alternative $\boldsymbol{x}$ follows from a Bernoulli distribution with a predictive distribution

$$\Pr(y=+1|\boldsymbol{x}, s) = \int \Pr(y=+1|\boldsymbol{x}, \boldsymbol{w}) \Pr(\boldsymbol{w}|s) \mathrm{d}\boldsymbol{w} = \int \sigma(\boldsymbol{w}^T \boldsymbol{x}) p(\boldsymbol{w}|s) \mathrm{d}\boldsymbol{w}. \tag{7}$$

We have

$$\begin{aligned} \mathbb{E}[V^{N+1}(T(s, \boldsymbol{x}, y))] &= \Pr(y=+1|\boldsymbol{x}, s) V^N(T(s, \boldsymbol{x}, +1)) + \Pr(y=-1|\boldsymbol{x}, s) V^N(T(s, \boldsymbol{x}, -1)) \\ &= \Pr(y=+1|\boldsymbol{x}, s) \cdot \max_{\boldsymbol{x}'} \Pr(y=+1|\boldsymbol{x}', T(s, \boldsymbol{x}, +1)) \\ &\quad + \Pr(y=-1|\boldsymbol{x}, s) \cdot \max_{\boldsymbol{x}'} \Pr(y=+1|\boldsymbol{x}', T(s, \boldsymbol{x}, -1)). \end{aligned}$$

The knowledge gradient policy suggests at each time n selecting the alternative that maximizes $\nu_{\boldsymbol{x}}^{\mathrm{KG}}(s^{n-1})$ where ties are broken randomly. The same optimization procedure as in recursive Bayesian logistic regression needs to be conducted for calculating the transition functions $T(s, \boldsymbol{x}, \cdot)$.

The predictive distribution $\int \sigma(\boldsymbol{w}^T \boldsymbol{x}) p(\boldsymbol{w}|s) \mathrm{d}\boldsymbol{w}$ cannot be evaluated exactly using the logistic function in the role of the sigmoid σ. An approximation procedure is deployed as follows. Denoting $a = \boldsymbol{w}^T \boldsymbol{x}$ and $\delta(\cdot)$ as the Dirac delta function, we have $\sigma(\boldsymbol{w}^T \boldsymbol{x}) = \int \delta(a - \boldsymbol{w}^T \boldsymbol{x}) \sigma(a) \mathrm{d}a$. Hence

$$\int \sigma(\boldsymbol{w}^T \boldsymbol{x}) p(\boldsymbol{w}|s) \mathrm{d}\boldsymbol{w} = \int \sigma(a) p(a) \mathrm{d}a,$$

where $p(a) = \int \delta(a - \boldsymbol{w}^T\boldsymbol{x})p(\boldsymbol{w}|s)\mathrm{d}\boldsymbol{w}$. Since $p(\boldsymbol{w}|s) = \mathcal{N}(\boldsymbol{m}, \boldsymbol{q}^{-1})$ is Gaussian, the marginal distribution $p(a)$ is also Gaussian. We can evaluate $p(a)$ by calculating the mean and covariance of this distribution [19]. We have

$$\begin{aligned}\mu_a &= \mathbb{E}[a] = \int p(a)a \ \mathrm{d}a = \int p(\boldsymbol{w}|s)\boldsymbol{w}^T\boldsymbol{x} \ \mathrm{d}\boldsymbol{w} = \hat{\boldsymbol{w}}^T\boldsymbol{x}, \\ \sigma_a^2 &= \int p(\boldsymbol{w}|s)\big((\boldsymbol{w}^T\boldsymbol{x})^2 - (\boldsymbol{m}^T\boldsymbol{x})^2\big) \ \mathrm{d}\boldsymbol{w} = \sum_{j=1}^{d} q_j^{-1}x_j^2.\end{aligned}$$

Thus $\int \sigma(\boldsymbol{w}^T\boldsymbol{x})p(\boldsymbol{w}|s)\mathrm{d}\boldsymbol{w} = \int \sigma(a)p(a)\mathrm{d}a = \int \sigma(a)\mathcal{N}(a|\mu_a, \sigma_a^2)\mathrm{d}a$.

For a logistic function, in order to obtain the best approximation [21, 22], we approximate $\sigma(a)$ by $\Phi(\alpha a)$ with $\alpha = \pi/8$. Denoting $\kappa(\sigma^2) = (1 + \pi\sigma^2/8)^{-1/2}$, we have

$$\Pr(y = +1|\boldsymbol{x}, s) = \int \sigma(\boldsymbol{w}^T\boldsymbol{x})p(\boldsymbol{w}|s)\mathrm{d}\boldsymbol{w} \approx \sigma(\kappa(\sigma^2)\mu).$$

We summarize the decision rule of the knowledge gradient policy at each iteration in Algorithm 2.

Algorithm 2: Knowledge Gradient Policy for Logistic Belief Model

input : m_j, q_j (Each weight w_j has an independent prior $\mathcal{N}(m_j, q_j^{-1})$)

for $\boldsymbol{x}$ *in* $\mathcal{X}$ **do**

Let $\Psi(\boldsymbol{w}, y) = -\frac{1}{2}\sum_{j=1}^{d} q_i(w_i - m_i)^2 - \log(1 + \exp(-y\boldsymbol{w}^T\boldsymbol{x}))$

$\hat{\boldsymbol{w}}_+ = \arg\max_{\boldsymbol{w}} \Psi(\boldsymbol{w}, +1)$

$\hat{\boldsymbol{w}}_- = \arg\max_{\boldsymbol{w}} \Psi(\boldsymbol{w}, -1)$

$\mu = \boldsymbol{m}^T\boldsymbol{x}$

$\sigma^2 = \sum_{j=1}^{d} q_j^{-1}x_j^2$

Define $\mu_+(\boldsymbol{x}') = \hat{\boldsymbol{w}}_+^T\boldsymbol{x}'$, $\mu_-(\boldsymbol{x}') = \hat{\boldsymbol{w}}_-^T\boldsymbol{x}'$

Define $\sigma_\pm^2(\boldsymbol{x}') = \sum_{j=1}^{d}\Big(q_j + \sigma(\hat{\boldsymbol{w}}_\pm^T\boldsymbol{x}')(1 - \sigma(\hat{\boldsymbol{w}}_\pm^T\boldsymbol{x}')x_j^2\Big)(x_j')^2$

$\tilde{\nu}_{\boldsymbol{x}} =$

$\sigma(\kappa(\sigma^2)\mu)\cdot\max_{\boldsymbol{x}'}\sigma\big(\kappa(\sigma_+^2(\boldsymbol{x}'))\mu_+(\boldsymbol{x}')\big) + \big(1 - \sigma(\kappa(\sigma^2)\mu)\big)\cdot\max_{\boldsymbol{x}'}\sigma\big(\kappa(\sigma_-^2(\boldsymbol{x}'))\mu_-(\boldsymbol{x}')\big)$

end

$\boldsymbol{x}^{\mathrm{KG}} = \arg\max_{\boldsymbol{x}} \tilde{\nu}_x$

We close this section by presenting the following finite-time bound on the MSE of the estimated weight with the proof in the supplement. Without loss of generality, we assume $\|\boldsymbol{x}\|_2 \leq 1$, $\forall \boldsymbol{x} \in \mathcal{X}$.

Theorem 1. *Let $\mathcal{D}^n$ be the n measurements produced by the KG policy and $\boldsymbol{w}^n = \arg\max_{\boldsymbol{w}} \Psi(\boldsymbol{w}|\boldsymbol{m}^0, \Sigma^0)$ with the prior distribution $Pr(\boldsymbol{w}^*) = \mathcal{N}(\boldsymbol{w}^*|\boldsymbol{m}^0, \boldsymbol{\Sigma}^0)$. Then with probability $P_d(M)$, the expected error of $\boldsymbol{w}^n$ is bounded as*

$$\mathbb{E}_{\boldsymbol{y}\sim\mathcal{B}(\mathcal{D}^n, \boldsymbol{w}^*)}||\boldsymbol{w}^n - \boldsymbol{w}^*||_2 \leq \frac{C_{min} + \lambda_{min}\big(\boldsymbol{\Sigma}^{-1}\big)}{2},$$

where the distribution $\mathcal{B}(\mathcal{D}^n, \boldsymbol{w}^*)$ is the vector Bernoulli distribution $Pr(y^i = +1) = \sigma(\boldsymbol{w}^{*T}\boldsymbol{x}^i)$, $P_d(M)$ is the probability of a d-dimension standard normal random variable appears in the ball with radius $M = \frac{1}{8}\frac{\lambda_{min}^2}{\sqrt{\lambda_{max}}}$ and $C_{min} = \lambda_{min}\Big(\frac{1}{n}\sum_{i=1}\sigma(\boldsymbol{w}^{*T}\boldsymbol{x}^i)\big(1 - \sigma(\boldsymbol{w}^{*T}\boldsymbol{x}^i)\big)\boldsymbol{x}^i(\boldsymbol{x}^i)^T\Big)$.

In the special case where $\boldsymbol{\Sigma}^0 = \lambda^{-1}\boldsymbol{I}$, we have $\lambda_{max} = \lambda_{min} = \lambda$ and $M = \frac{\lambda^{3/2}}{8}$. The bound holds with higher probability $P_d(M)$ with Iarger λ. This is natural since a larger λ represents a normal distribution with narrower bandwidth, resulting in a more concentrated $\boldsymbol{w}^*$ around $\boldsymbol{m}^0$.

4 Experimental results

We evaluate the proposed method on both synthetic datasets and the UCI machine learning repository [23] which includes classification problems drawn from settings including fertility, glass identification, blood transfusion, survival, breast cancer (wpbc), planning relax and climate model failure. We first analyze the behavior of the KG policy and then compare it to state-of-the-art active learning algorithms. On synthetic datasets, we randomly generate a set of M d-dimensional alternatives $\boldsymbol{x}$ from [-3,3]. We conduct experiments in a Bayesian fashion where at each run we sample a true $d+1$-dimensional weight vector $\boldsymbol{w}^*$ from the prior distribution $w_i^* \sim \mathcal{N}(0, \lambda)$. The $+1$ label for each alternative $\boldsymbol{x}$ is generated with probability $\sigma(w_0^* + \sum_{j=1}^{d} w_d^* x_d)$. For each UCI dataset, we use all the data points as the set of alternatives with their original attributes. We then simulate their labels using a weight vector $\boldsymbol{w}^*$. This weight vector could have been chosen arbitrarily, but it was in fact a perturbed version of the weight vector trained through logistic regression on the original dataset. All the policies start with the same one randomly selected example per class.

4.1 Behavior of the KG policy

To better understand the behavior of the KG policy, Fig. 1 shows the snapshot of the KG policy at each iteration on a 2-dimensional synthetic dataset ($M = 200$) in one run. The scatter plots show the KG values with both the color and the size of the point reflecting the KG value of the corresponding alternative. The star denotes the true alternative with the largest response. The red square is the alternative with the largest KG value. The pink circle is the implementation decision that maximizes the response under current estimation of $\boldsymbol{w}^*$ if the budget is exhausted after that iteration.

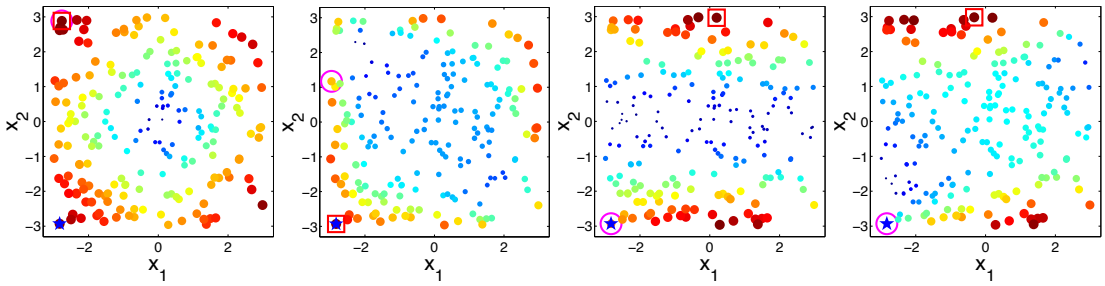

Figure 1: The scatter plots illustrate the KG values at 1-4 iterations from left to right with both the color and the size reflecting the magnitude. The star, the red square and pink circle indicate the true best alternative, the alternative to be selected and the implementation decision, respectively.

It can be seen from the figure that the KG policy finds the true best alternative after only three measurements, reaching out to different alternatives to improve its estimates. We can infer from Fig. 1 that the KG policy tends to choose alternatives near the boundary of the region. This criterion is natural since in order to find the true maximum, we need to get enough

information about $\boldsymbol{w}^*$ and estimate well the probability of points near the true maximum which appears near the boundary. On the other hand, in a logistic model with labeling noise, a data $\boldsymbol{x}$ with small $\boldsymbol{x}^T\boldsymbol{x}$ inherently brings little information as pointed out in [24]. For an extreme example, when $\boldsymbol{x} = \mathbf{0}$ the label is always completely random for any $\boldsymbol{w}$ since $\Pr(y = +1|\boldsymbol{w}, \mathbf{0}) \equiv 0.5$. This is an issue when perfect classification is not achievable. So it is essential to label a data with larger $\boldsymbol{x}^T\boldsymbol{x}$ that has the most potential to enhance its confidence non-randomly.

Also depicted in Fig. 2 is the absolute class distribution error of each alternative, which is the absolute difference between the predictive probability of class +1 under current estimate and the true probability after 6 iterations. We see that the probability at the true maximum is well approximated, while moderate error in the estimate is located away from this region of interest. We also provide the analysis on a 3-dimensional dataset in the supplement.

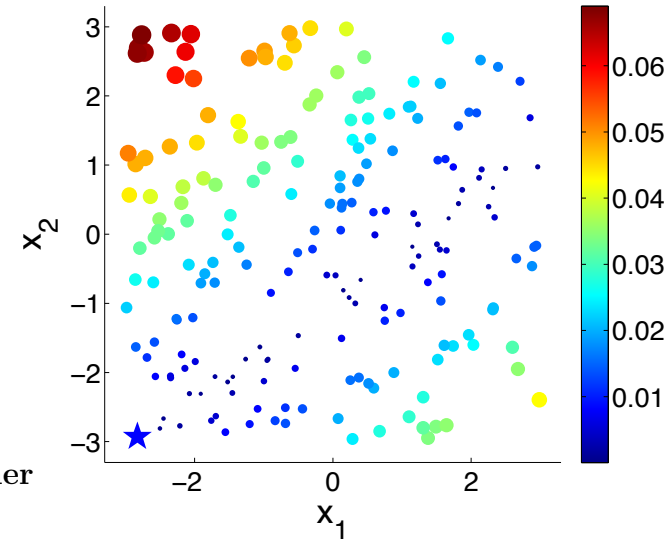

Figure 2: Absolute error.

4.2 Comparison with other policies

Recall that our goal is to maximize the expected response of the implementation decision. We define the Opportunity Cost (OC) metric as the expected response of the implementation decision $\boldsymbol{x}^{N+1} := \arg\max_{\boldsymbol{x}} \Pr(y = +1|\boldsymbol{x}, \boldsymbol{w}^N)$ compared to the true maximal response under weight $\boldsymbol{w}^*$:

$$\text{OC} := \max_{\boldsymbol{x}\in\mathcal{X}} \Pr(y = +1|\boldsymbol{x}, \boldsymbol{w}^*) - \Pr(y = +1|\boldsymbol{x}^{N+1}, \boldsymbol{w}^*).$$

Note that the opportunity cost is always non-negative and the smaller the better. To make a fair comparison, on each run, all the time-N labels of all the alternatives are randomly pre-generated according to the weight vector $\boldsymbol{w}^*$ and shared across all the competing policies.

Since there is no policy directly solving the same sequential response maximizing problem under a logistic model, considering the relationship with active learning as described in Section 1, we compare with the following state-of-the-art active learning policies compatible with logistic regression: Random sampling (Random), a myopic method that selects the most uncertain instance each step (MostUncertain), Fisher information (Fisher) [25], the batch-mode active learning via error bound minimization (Logistic Bound) [26] and discriminative batch-mode active learning (Disc) [27] with batch size equal to 1. All the state transitions are based on recursive Bayesian logistic regression while different policies provides different rules for labeling decisions at each iteration. The experimental results are shown in figure 3. In all the figures, the x-axis denotes the number of measured alternatives and the y-axis represents the averaged opportunity cost over 100 runs.

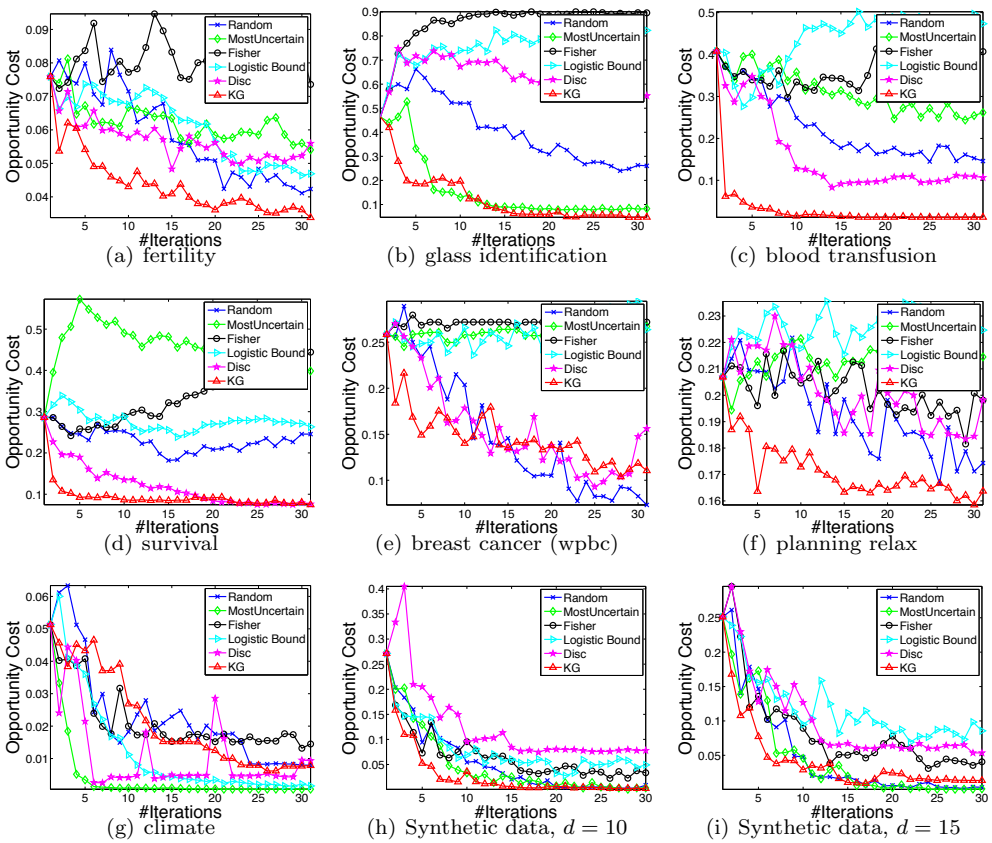

(a) fertility (b) glass identification (c) blood transfusion

(d) survival (e) breast cancer (wpbc) (f) planning relax

(g) climate (h) Synthetic data, $d = 10$ (i) Synthetic data, $d = 15$

Figure 3: Opportunity cost on UCI and synthetic datasets.

It is demonstrated in FIG. 3 that KG outperforms the other policies significantly in most cases, especially in early iterations. MostUncertain, Fisher and Logistic Bound perform well on some datasets while badly on others. Disc and Random yield relatively stable and satisfiable performance. A possible explanation is that the goal of active leaning is to learn a classifier which accurately predicts the labels of new examples so their criteria are not directly related to maximize the response aside from the intent to learn the prediction. After enough iterations when active learning methods presumably have the ability to achieve a good estimator of $\boldsymbol{w}^*$, their performance will be enhanced. However, in the case when an experiment is expensive and only a small budget is allowed, the KG policy, which is designed specifically to maximize the response, is preferred.

5 Conclusion

In this paper, we consider binary classification problems where we have to run expensive experiments, forcing us to learn the most from each experiment. The goal is to learn the classification model as quickly as possible to identify the alternative with the highest response. We develop a knowledge gradient policy using a logistic regression belief model, for which we developed an approximation method to overcome computational challenges in finding the knowledge gradient. We provide a finite-time analysis on the estimated error, and report the results of a series of experiments that demonstrate its efficiency.

References

[1] M. H. DeGroot. *Optimal Statistical Decisions*. McGraw-Hill, 1970.

[2] G. B. Wetherill and K. D. Glazebrook. *Sequential Methods in Statistics*. Chapman and Hall, 1986.

[3] D. C. Montgomery. *Design and Analysis of Experiments*. John Wiley and Sons, 2008.

[4] Andrew I Schein and Lyle H Ungar. Active learning for logistic regression: an evaluation. *Machine Learning*, 68(3):235–265, 2007.

[5] Simon Tong and Daphne Koller. Support vector machine active learning with applications to text classification. *The Journal of Machine Learning Research*, 2:45–66, 2002.

[6] Yoav Freund, H Sebastian Seung, Eli Shamir, and Naftali Tishby. Selective sampling using the query by committee algorithm. *Machine learning*, 28(2-3):133–168, 1997.

[7] Burr Settles. Active learning literature survey. *University of Wisconsin, Madison*, 52(55-66):11, 2010.

[8] Peter Auer, Nicolò Cesa-Bianchi, and Paul Fischer. Finite-time analysis of the multi-armed bandit problem. *Machine learning*, 47(2-3):235–256, 2002.

[9] Sébastien Bubeck and Nicolò Cesa-Bianchi. Regret analysis of stochastic and nonstochastic multi-armed bandit problems. *arXiv preprint arXiv:1204.5721*, 2012.

[10] Ilya O Ryzhov, Warren B Powell, and Peter I Frazier. The knowledge gradient algorithm for a general class of online learning problems. *Operations Research*, 60(1):180–195, 2012.

[11] Donghai He, Stephen E Chick, and Chun-Hung Chen. Opportunity cost and OCBA selection procedures in ordinal optimization for a fixed number of alternative systems. *Systems, Man, and Cybernetics, Part C: Applications and Reviews, IEEE Transactions on*, 37(5):951–961, 2007.

[12] Stephen E Chick. New two-stage and sequential procedures for selecting the best simulated system. *Operations Research*, 49(5):732–743, 2001.

[13] Warren B Powell and Ilya O Ryzhov. *Optimal learning*. John Wiley & Sons, 2012.

[14] Peter I Frazier, Warren B Powell, and Savas Dayanik. A knowledge-gradient policy for sequential information collection. *SIAM Journal on Control and Optimization*, 47(5):2410–2439, 2008.

[15] Martijn RK Mes, Warren B Powell, and Peter I Frazier. Hierarchical knowledge gradient for sequential sampling. *The Journal of Machine Learning Research*, 12:2931–2974, 2011.

[16] Diana M Negoescu, Peter I Frazier, and Warren B Powell. The knowledge-gradient algorithm for sequencing experiments in drug discovery. *INFORMS Journal on Computing*, 23(3):346–363, 2011.

[17] Yingfei Wang, Kristofer G Reyes, Keith A Brown, Chad A Mirkin, and Warren B Powell. Nested-batch-mode learning and stochastic optimization with an application to sequential multistage testing in materials science. *SIAM Journal on Scientific Computing*, 37(3):B361–B381, 2015.

[18] David W Hosmer Jr and Stanley Lemeshow. *Applied logistic regression*. John Wiley & Sons, 2004.

[19] Christopher M Bishop et al. *Pattern recognition and machine learning*, volume 4. springer New York, 2006.

[20] Stephen J Wright and Jorge Nocedal. *Numerical optimization*, volume 2. Springer New York, 1999.

[21] David Barber and Christopher M Bishop. Ensemble learning for multi-layer networks. *Advances in neural information processing systems*, pages 395–401, 1998.

[22] David J Spiegelhalter and Steffen L Lauritzen. Sequential updating of conditional probabilities on directed graphical structures. *Networks*, 20(5):579–605, 1990.

[23] M. Lichman. UCI machine learning repository, 2013.

[24] Tong Zhang and F Oles. The value of unlabeled data for classification problems. In *Proceedings of the Seventeenth International Conference on Machine Learning,(Langley, P., ed.)*, pages 1191–1198. Citeseer, 2000.

[25] Steven CH Hoi, Rong Jin, Jianke Zhu, and Michael R Lyu. Batch mode active learning and its application to medical image classification. In *Proceedings of the 23rd international conference on Machine learning*, pages 417–424. ACM, 2006.

[26] Quanquan Gu, Tong Zhang, and Jiawei Han. Batch-mode active learning via error bound minimization. *Urbana*, 51:61801, 2014.

[27] Yuhong Guo and Dale Schuurmans. Discriminative batch mode active learning. In *Advances in neural information processing systems*, pages 593–600, 2008.